

New developments in event generator tuning techniques

Andy Buckley*University of Edinburgh**E-mail:* andy.buckley@ed.ac.uk**Hendrik Hoeth***IPPP Durham**E-mail:* hendrik.hoeth@cern.ch**Heiko Lacker***Humboldt-University Berlin**E-mail:* lacker@physik.hu-berlin.de**James Monk****UCL**E-mail:* jmonk@hep.ucl.ac.uk**Holger Schulz***Humboldt-University Berlin**E-mail:* hschulz@physik.hu-berlin.de**Jan Eike von Seggern***Humboldt-University Berlin**E-mail:* vseggern@physik.hu-berlin.de

Data analyses in hadron collider physics depend on background simulations performed by Monte Carlo (MC) event generators. However, calculational limitations and non-perturbative effects require approximate models with adjustable parameters. In fact, we need to simultaneously tune many phenomenological parameters in a high-dimensional parameter-space in order to make the MC generator predictions fit the data. It is desirable to achieve this goal without spending too much time or computing resources iterating parameter settings and comparing the same set of plots over and over again. We present extensions and improvements to the MC tuning system, Professor, which addresses the aforementioned problems by constructing a fast analytic model of a MC generator which can then be easily fitted to data. Using this procedure it is for the first time possible to get a robust estimate of the uncertainty of generator tunings. Furthermore, we can use these uncertainty estimates to study the effect of new (pseudo-) data on the quality of tunings and therefore decide if a measurement is worthwhile in the prospect of generator tuning. The potential of the Professor method outside the MC tuning area is presented as well.

*13th International Workshop on Advanced Computing and Analysis Techniques in Physics Research,
ACAT2010
February 22-27, 2010
Jaipur, India*

1. MC tuning with Professor

The “Professor” approach to MC tuning constitutes both a numerical method and a suite of tools which implement it. Fundamentally, Professor attempts to parameterise expensive functions – the bin values in a set of MC observables – by least-squares fitting of the parameterisation coefficients. The least-squares minimisation is made more approachable by use of the pseudoinverse method, implemented via a matrix singular value decomposition. Armed with a fast analytic model of how every bin of a large set of observables will respond to variations of the generator parameters, numerical optimisation of the generator’s fit to reference data may be efficiently computed. A detailed description may be found in reference [1].

The benefit of this approach is clear for more than 2 parameters: Professor requires as input the values of observables for a moderately large number of MC runs distributed suitably in the generator parameter space, each point in the space perhaps requiring $\mathcal{O}(48)$ CPU hours to complete. A serial optimisation approach such as Markov chain sampling would hence require thousands of CPU days to have a chance of converging, if the generator itself is not batch-parallelised. Professor, given sufficiently large batch computing resources, can trivially parallelise the generation of the input MC points for any generator and thereafter complete the parameterisation and fit optimisation in negligible time, allowing for scaling to higher numbers of tune parameters than could be attempted by methods which require iteration of the time-limiting step.

2. Qualitative error estimation in Professor

An important feature of the Professor method is that it has always allowed for *qualitative* assessment of the tune robustness. A per-bin parameterisation in p parameters will require a minimum number of MC runs, $N_{\min}^{(p)}$, for the least-squares pseudoinversion to be performed. For robustness it is advisable to oversample this minimum requirement by a factor $\mathcal{O}(3)$ (or more, especially for large p) such that a tune will in fact use $N \gg N_{\min}^{(p)}$ input runs. Additionally, since the parameterisation and optimisation steps are fast, we take the opportunity to make many such overconstrained tunes by in fact sampling an even greater number of runs, $N_{\text{sampled}} \gg N$. We can then randomly sample a large number of mostly-independent N -run tunes to obtain an ensemble of reasonable tunes – again, this step can be trivially batch-parallelised. The spread of this tune ensemble as projected on each parameter has been used in several Professor MC tunes as a heuristic for determining whether a parameter is well or poorly constrained, for detecting parameterisation problems, and for ensuring that the “maximum information” tune is typical of the ensemble. Several other checking methods, such as eigenvector line scans, are also used to ensure that the details of the tune, and particularly the generator parameterisation, are reliable.

To make a prediction of an observable for which there is no reference data (this may be an energy extrapolation, or simply an unmeasured feature at existing energies), the simplest approach is of course to run the generator with the obtained tune(s) and compute the observable. Using the many tunes resulting from the different run combinations will give a spread in the observable prediction, reflecting one part of tune uncertainty. In practice, since we can build a fast parameterisation of the

*Many thanks to James Monk who presented this material in our absence.




Figure 1: A scatter plot of χ^2 vs. parameter value for a set of Professor run combinations on three Pythia6 parameters. Implicitly, this projection of tune parameter vectors on to parameter axes gives a qualitative measure of whether or not a parameter is well-constrained: these parameters become increasingly ill-defined from left to right. The different markers represent different degrees of oversampling, with the star representing the maximum information run — the points are for the same run combinations in all three scatter plots.

generator behaviour (in fact many of them) on the unfitted observable as part of the main Professor tuning process, this offers a much faster turnaround than processing another large (perhaps *very* large) set of generator runs – with the proviso that the parameterisations are of course non-exact.

3. Sources of tune uncertainty

Before making this method more quantitative, we now consider the various sources of uncertainty in the procedure outlined so far. This will help us to understand which sources of uncertainty are computationally controllable and which will have to, for now, remain more nebulous. These main sources of tuning error are as follows:

1. Error on experimental reference data.
2. Statistical error on the MC at the anchor points from which the parameterisation is constructed.
3. Systematic limitations of the parameterising function to describe bin responses to parameter variations – pathological MC parameters with discontinuous or critical behaviour are particularly hard to generically parameterise, since a Jacobian transformation to a suitable meta-parameter is not always easily available.
4. Choice of run combinations to make the parameterisation.
5. Goodness of fit definition, including both the type of GoF measure and the choice and relative weighting of different data.
6. Reasonable minimiser scatter within the χ^2 valley containing the optimal tune point for a given parameterisation. Note that this cannot be completely disentangled from the role that error sources 1, 2 and 3 play in defining the χ^2 valley.
7. Limitations of the parameterisation(s) used to compute the parameterised MC value in extrapolated/unfitted observables. Of course, this error doesn't exist if the less convenient strategy of re-running the MC generator is used.

8. For completeness, we again highlight the systematic error associated with the discrete physics model being tuned: the total error is far from complete without considering more than one viable model. Within a given model there are also systematic uncertainties, some of which may be quantified, e.g. cross-section integration uncertainty and PDFs: the second of these is particularly quantifiable due to the existence of error or replica PDF sets, themselves expressing reasonable variations in PDF fitting.

Note that, for example, these sources of uncertainty such as the variation between members of the ensemble of run combinations are not unique errors introduced by the Professor approach: failing to test different run combinations does not *eliminate* the error associated with the choice of anchor runs used! A similar, but unquantifiable, error exists for any form of manual tuning.

4. Construction of tune uncertainty confidence belts

Our approach to quantifying the uncertainties from (combinations of) the sources listed in the previous section is to construct central confidence belts for observable bins from the various ensembles of tune results, parameterisations, etc. we have described. Explicitly, given a large number of reasonable and equivalent predictions for an observable bin value, we construct a band of given P -value as being the region containing fraction P of predictions, with equal fractions above and below.

There are many ways to construct such ensembles – for the purposes of this study we identify three:

Combination error: The ensemble from which we construct the confidence belt is simply the ensemble of predictions from different run combinations. This will hence represent the variation due to error sources 1, 2, and 4 – the other sources of uncertainty exist, but are not quantified by this approach¹. This approach requires that correlations between different run combinations are small, which is ensured by the $N_{\text{sampled}} \gg N$ requirement. If parameterisation (as opposed to explicit MC runs) is used for the translation of this tune ensemble into bin value predictions, then source 7 also applies. This can be included into the band construction by using many parameterisations, again constructed from run combinations. Different parameterisations should be used for minimisation and prediction to avoid reinforcing parameterisation systematics.

Statistical error: The obvious failure of the combination error approach is that error source 6 – the measure of reasonable tune variation within the χ^2 valley – is left unquantified. Hence it will be no surprise that our “statistical” error band is constructed explicitly from an ensemble of samples from this valley. This is obtained in the simplest case by only using the maximum information Professor tune – that which is constructed from all the available MC runs. The covariance matrix of the parameters in the vicinity of the tune point is obtained – in principle directly from the parameterisation, in practice from the minimiser – and used to define a rotated hyper-Gaussian probability distribution in the parameter space. Sampling parameter points from

¹In the results shown here, MC error has not been *explicitly* propagated into Professor’s fit measure, due to extrapolation problems: it enters implicitly via the statistical scatter of MC samples. This has been remedied in the latest development versions of Professor, which use the median MC error of the sampled anchor points to avoid the instabilities.




Figure 2: Two dimensional illustration of the parameter point sampling used for the statistical uncertainty estimate. We exploit the covariance matrix returned by the minimiser for a Gaussian sampling from the corresponding p -dimensional hyper-ellipsoid. The σ_i ($i = 1 \dots p$) are the eigenvalues obtained from an eigen-decomposition of the covariance matrix.

this distribution gives us another ensemble of tune points and, as for the combination error, they can be mapped into observable predictions either by direct MC runs or by using one or many constructed parameterisations. In practice, we take advantage of the combined potential of the hyper-Gaussian sampling and MC parameterisation to build a confidence belt from $\mathcal{O}(10,000)$ samples. The quantified sources of error are hence 1, 2, 6, and 7.

These approaches to error band building will be used in the next section to construct sample error bands for underlying event predictions. Although not currently implemented in Professor, we also highlight the most complete form of quantifiable error band within the Professor approach:

Combined error: This extension is an obvious fusion of the above two ensemble/band constructions: as in the combination error, we construct an ensemble of points from run combinations, then for each run-combination point we construct a statistical ensemble. The combined ensemble of tune ensembles, and a variety of parameterisations to transform them into predictions will lead to error bands quantifying error sources 1, 2, 4, 6, and 7.

The remaining error sources are 3, 5 and 8: limitations of the parameterisation, the observable weights/goodness of fit definition, and – most importantly – the uncertainty due to different physics models. These missing systematics remain qualitative in this scheme, and reliable MC predictions should take care to include estimates of their influence, albeit in a more ad hoc fashion than for the more statistically-induced errors.

5. Results

We now briefly present results using the first two confidence belt definitions presented above – ongoing work is addressing the “combined” belt definition and inclusion of PDF uncertainties.

Our exploration is based on tunes of the Jimmy [2] MC generator, which simulates multiple parton interactions (MPI) for Herwig 6 [3], because it has only two relevant parameters, PTJIM and JMRAD(73)² – the frugality with parameters makes Jimmy an ideal “toy model” testbed, while remaining phenomenologically relevant. As a Jimmy-like MPI model is ruled out by Tevatron data [4], we fix a dependence of PTJIM on the centre of mass energy with the same ansatz as used in Pythia 6 [5]:

$$\text{PTJIM} = \text{PTJIM}_{1800} \cdot \left(\frac{\sqrt{s}}{1800 \text{ GeV}} \right)^{0.274}, \quad (5.1)$$

²We treat the inverse radius-squared of the protons, JMRAD(73), to be identical to that of the anti-protons, JMRAD(91), and ignore the interplay with ISR parameters in Herwig 6 itself.


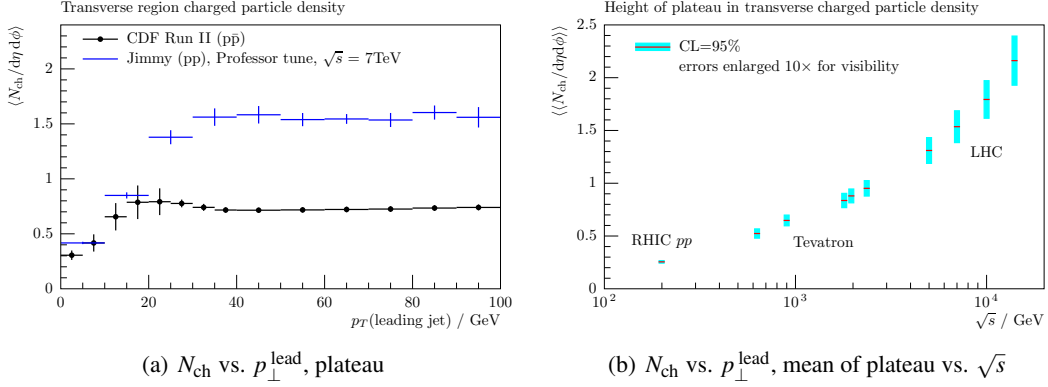


Figure 3: “Statistical” and “combination” error bands for transverse N_{ch} flow at 7 TeV before and after adding 1M events of pseudodata of this observable (black markers) to the tuning. The error bands are calculated from the central 95 (68)% of the binvalues of an ensemble of 10000 histograms each.

where PTJIM_{1800} is the value of PTJIM at the reference scale $\sqrt{s} = 1800$ GeV and is the p_{\perp}^{min} parameter actually used in the tuning process. Furthermore, we use the MRST LO* PDF set [6] and use Tevatron data from CDF [7–9] and DØ [10] as a tuning reference. A more complete tune would include the exponent of the p_{\perp}^{min} energy dependence, but for toy-study purposes we here fix it to a value consistent with known energy extrapolation fits [1, 11].

In Figures 3(a) and 3(c), the statistical and combination error band definitions are shown for a 7 TeV underlying event observable – N_{ch} flow transverse to the leading track (track with the largest transverse momentum in an event), as a function of leading track p_{\perp} – computed in Rivet [12], based on the fit to Tevatron reference data. The error due to variation of run combinations (error source 4) is notably somewhat larger than the scatter of points in the χ^2 valley (error source 6), indicating that in the Jimmy model the parameters have a strong influence on this observable, and are hence well-constrained. In Figures 3(b) and 3(d), similar band constructions are shown, but in this case the same Rivet analysis has been used to simulate the effect of adding 1M events of UE data at 7 TeV into the fit: the size of both error bands is reduced, as expected.


 (a) N_{ch} vs. p_{\perp}^{lead} , plateau

 (b) N_{ch} vs. p_{\perp}^{lead} , mean of plateau vs. \sqrt{s}

5.1 Effect of extrapolation

Finally, we consider systematically how error bands constructed in this way behave as extrapolations are taken further from the region of constraining data. In this case, since the computational requirements are significant, we only consider the “statistical” error band.

We use the transverse region N_{ch} density UE observable, evaluated at a range of ten centre of mass energies, \sqrt{s}_i , between 200 GeV and 14 TeV. We then apply the following procedure,

- Construct the maximum-information parameterisation of the generator response for the N_{ch} UE observable, $f(\sqrt{s}_i)$
- Produce an ensemble of 10000 histograms, H_i , of the observable shown in Figure 4(a) (blue line) using the corresponding $f(\sqrt{s}_i)$ and points sampled using the procedure illustrated in Figure 2
- Calculate the mean height of the N_{ch} plateau, M_i , for each of the H_i
- Construct a 95% central confidence belt, $CL(\sqrt{s}_i)$ from the M_i

In Figure 4(b) the $CL(\sqrt{s}_i)$ are drawn. We observe a very tight confidence belt for the energy region of the Tevatron experiments, while the confidence belt becomes wider for extrapolation to LHC energies. The definitions of the plateau regions used can be found in Table 1.

It is notable that these bands are narrow – sufficiently so that they have been visually inflated by a factor of 10 in the figure. While this reflects good stability in the tuning system, it is probably an underestimate of the true model uncertainty. A more complete study will include the exponent of energy extrapolation in equation (5.1) in the tune, since this may be a dominant effect in the errors for this particular observable and its inclusion will give more freedom for various features of the model to balance against each other: the toy tuning of the energy-dependent Jimmy model shown here is probably too restrictive to accurately represent the full range of variation allowed by the model, but serves as an indication of the extent to which such studies can be systematised.

\sqrt{s}/TeV	0.2	0.63	0.9	1.8	1.96	2.36	5.0	7.0	10.0	14.0
$p_{\perp}^{\text{lead, min}}/\text{GeV}$	10	30	30	30	40	40	40	40	40	40
$p_{\perp}^{\text{lead, max}}/\text{GeV}$	30	70	80	80	90	110	110	120	150	160

Table 1: Definition of plateau regions (p_{\perp}^{lead}) used in the extrapolation study.

6. Interactive parameterisation explorer

The Professor package also contains the program `prof-I` that allows for interactive exploration of the effect of changing parameter values on the shape of observables. It makes use of the functionality of the WXPYthon [13] library to create a graphical user interface (GUI). The user can choose a Professor interpolation to explore and overlay experimental and Monte Carlo generator (pseudo-) data for two different observables.

The heart of this program are simple sliders (one for each parameter in the interpolation) that can be adjusted freely within certain values of the corresponding parameter. Exploring the parameter space is done simply by moving the sliders which results in real-time redrawing of the interpolation's prediction for the two observables studied.

By doing so it one immediately gets a prediction of the generators outcome at a certain parameter point, a task that by running the generator alone usually takes many hours or days. This program in principle enables the user to do manual tunings tool, however, the much greater value of this tool is that it helps to get a feeling for the effects of shifts in parameter space on certain observables. Also the decision which observables one should include in a tuning can be helped since the sensitivity of observables to parameters can be studied with the `prof-I` GUI.

7. Conclusions

We have catalogued a set of sources of uncertainty which either explicitly or implicitly contribute to any tune of a MC event generator, and presented a systematic approach to quantifying many of these uncertainties using the fast MC parameterisations and natural tune ensembles which arise from the Professor tuning approach. Example results have been shown, which exhibit some expected behaviours, such as the shrinking of error bands on adding new reference data in new areas of parameter space and the blow-up of error bands as predictions venture further into unconstrained regions.

Several things should be emphasised: first and most important is that this approach does not catch all sources of error. We have presented results from two definitions and have proposed a third, more comprehensive measure, but still variations such as PDF errors and discrete model variations need to be included. However, with the increased usage of systematic tuning methods, variations between models in UE observables are not as substantial as once they were – statistical errors are a non-negligible factor in assessing the reliability of phenomenologically-based MC predictions.

The approach taken here has many obvious parallels in the world of PDF errors, with our approach having a good deal of overlap with the MC replica set approach of the NNPDF collaboration [14] as contrasted with the eigenset approach of the CTEQ and MRST/MSTW collaborations.




Figure 4: A screenshot of the interactive parameter explorer of the Professor package, `prof-I`. The blue line represents the interpolation’s prediction for points in parameter space that are adjusted by the sliders at the bottom of the GUI: moving the sliders results in real-time redrawing.

While replica sets have the advantage of a more direct statistical uncertainty interpretation (although we do not have the option of the parameterisation-freedom exhibited by the NNPDF use of neural networks), there is the pragmatic issue that $\mathcal{O}(10)$ representative error tunes would be more usable than $\mathcal{O}(10000)$ equivalent tunes. Work to add “eigen-tune” calculation to Professor has begun, and we hope to soon provide statistically-driven rivals to the Perugia-soft/hard tunes for systematic sensitivity studies.

Acknowledgements

We would like to thank James Monk for presenting this material in our absence. Our particular thanks goes to Luigi del Debbio and Richard Ball for discussions about statistical coverage and MC error estimation. The Professor collaboration acknowledges support from the EU MCnet Marie Curie Research Training Network (funded under Framework Programme 6 contract MRTN-CT-2006-035606) for financial support and for many useful discussions and collaborations with its members. A. Buckley additionally acknowledges support from the Scottish Universities Physics Alliance (SUPA); H. Schulz acknowledges the support of the German Research Foundation (DFG).

References

- [1] A. Buckley, H. Hoeth, H. Lacker, H. Schulz and J. E. von Seggern, “Systematic event generator tuning for the LHC,” *Eur. Phys. J. C* **65** (2010) 331 [arXiv:0907.2973 [hep-ph]].
- [2] J. M. Butterworth and M. H. Seymour, “JIMMY4: Multiparton Interactions in Herwig for the LHC”, 2004, [<http://projects.hepforge.org/jimmy>].
- [3] G. Corcella *et al.*, “HERWIG 6.5 release note,” [arXiv:hep-ph/0210213].
- [4] M. Bahr, J. M. Butterworth and M. H. Seymour, “The Underlying Event and the Total Cross Section from Tevatron to the LHC,” *JHEP* **0901** (2009) 065 [arXiv:0806.2949 [hep-ph]].
- [5] T. Sjostrand, S. Mrenna and P. Z. Skands, “PYTHIA 6.4 Physics and Manual,” *JHEP* **0605** (2006) 026 [arXiv:hep-ph/0603175].
- [6] A. Sherstnev and R. S. Thorne, “Parton Distributions for LO Generators,” *Eur. Phys. J. C* **55** (2008) 553 [arXiv:0711.2473 [hep-ph]].
- [7] T. Affolder and others (CDF), “Charged jet evolution and the underlying event in $p\bar{p}$ collisions at 1.8 TeV” *Phys. Rev.*, D65, 2002.
- [8] D. E. Acosta *et al.* [CDF Collaboration], “The underlying event in hard interactions at the Tevatron $p\bar{p}$ collider,” *Phys. Rev. D* **70** (2004) 072002 [arXiv:hep-ex/0404004].
- [9] The CDF Collaboration, “Studying the Underlying Event in Drell-Yan and High Transverse Momentum Jet Production at the Tevatron,” arXiv:1003.3146.
- [10] V. M. Abazov *et al.* [D0 Collaboration], “Measurement of dijet azimuthal decorrelations at central rapidities in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV,” *Phys. Rev. Lett.* **94** (2005) 221801 [arXiv:hep-ex/0409040].
- [11] P. Z. Skands, “The Perugia Tunes,” arXiv:0905.3418 [hep-ph].
- [12] B. M. Waugh, H. Jung, A. Buckley, L. Lonnblad, J. M. Butterworth and E. Nurse, “HZTool and Rivet: Toolkit and framework for the comparison of simulated final states and data at colliders,” arXiv:hep-ph/0605034.
- [13] R. Dunn, H. Pasanen, “WXPython: A GUI toolkit for the Python programming language.”, <http://www.wxpython.org/>.
- [14] R. D. Ball *et al.* [NNPDF Collaboration], “A determination of parton distributions with faithful uncertainty estimation,” *Nucl. Phys. B* **809** (2009) 1 [Erratum-ibid. B **816** (2009) 293] [arXiv:0808.1231 [hep-ph]].